

Comparing Properties of Massively Multiplayer Online Worlds and the Internet of Things [v-0.9]

Kim J.L. Nevelsteen, Theo Kanter, Rahim Rahmani

Abstract

With the rise of the Internet of Things (IoT), this means recognizing the need for architectures to handle billions of devices and their interactions. A virtual world engine at the massively multiplayer scale is a massively multiplayer online world (MMOW); one thing virtual world engines realized when going into the scale of MMOs, is the cost of maintaining a potentially quadratic number of interactions between a massive number of objects, laid out in a spatial dimension. Research into IoT was fueled by research in wireless sensor networks, but rather than start from a device perspective, this article looks at how architectures deal with interacting entities at large scale. The domain of MMOWs is examined for properties that are affected by scale. Thereafter the domain of IoT is evaluated to see if each of those properties are found and how each is handled. By comparing the current state of the art of MMOWs and IoT, with respect to scalability, the problem of scaling IoT is explicated, as well as the problem of incorporating an MMOW with IoT into a pervasive platform. Three case studies of a MMOW interfacing with IoT are presented in closing.

1 Introduction

A virtual world at the massively multiplayer scale is a Massively Multiplayer Online World (MMOW) [18]. One thing virtual world engines realized when going into the scale of MMOs, is the cost of maintaining a potentially quadratic number of interactions between a massive number of objects, laid out in a spatial dimension [23, 6]. Yahyavi and Kemme [23] explicitly state that the architectures they focus on for MMO games are also applicable to other distributed systems *e.g.*, technology-sustained pervasive games [17].

With the rise of the Internet of Things (IoT) [4, 19, 1], this means recognizing the need for architectures to handle billions of devices and their interactions. Endeavors are already underway in an attempt to create an IoT platform, but a “solution that addresses all the aspects required by the IoT is yet to be designed” [19]. Miorandi et al. [15] summarize a number of research initiatives happening worldwide *e.g.*, HYDRA allowing developers to incorporate heterogeneous devices and IoT-A concentrating on interoperability. Devices in IoT (*e.g.*, RFID) allow for the mapping of the physical world into the virtual world [1] *i.e.*, using ‘non-standard input devices’ to blend the virtual and the physical [17] into a pervasive system [20].

Research into IoT was fueled by research in Wireless Sensor and Actuator Networks (WSANs) [4], but rather than start from a device perspective, in Section 2, the domain of MMOWs is examined to see how architectures deal with scalability *i.e.*, properties that are affected by scale are gathered from the domain of MMOWs. In Section 3, it is then evaluated if each of those properties is found in the domain of IoT and how each property is handled. By comparing the current state of the art of MMOWs and IoT, with respect to scalability, Section 4 points to how research from one domain can possibly be exapted to the other domain and *vice versa*. In addendum, Section 5 provides three case studies as to how an MMOW can interface with IoT or how a MMOW can be a ‘mediator’ [20] for IoT. The article closes with conclusions summarized in Section 6.

2 Examining the Domain of MMOWs

Since virtual world engines (implementing MMOWs) have long been dealing with the cost of maintaining a massive number of objects, issues are gathered from the domain of MMOWs that are affected by scale. Issues are categorized according to properties presented in the ISO 25010:2011 standard [8].

An issue, relative to MMOWs, that is currently under research is that of **scalability**. “Scalability can be achieved either by:(1) increasing the resources or by (2) reducing the consumption” [23]. Adding multiple servers to distribute the load of handling interactions is designed to increase the amount of resources. To alleviate the scalability problem, a server cluster can be used instead of the single server *i.e.*, a co-located cluster of servers that collectively act as a centralized unit to serve all clients. To resolve the scalability issue further, current MMO research is looking into pure peer-to-peer (P2P) solutions or a hybrid P2P server cluster combination. Server clusters can be formed in a distributed P2P system by assigning a ‘region controller’ [5] to supervise over the peers in the cluster. If multiple servers are used, two ways to partition computational space are regions and shards *i.e.*, regionalization and replication. Regionalization divides space into regions, with a different set of servers responsible for a different region, and replication means having multiple copies of the same space. With replicated shards there is no or minimal interaction between shards [23]. A hybrid is also possible *e.g.*, a shard divided into regions. To achieve scalability through a decrease of consumption, ‘interest management’ can be used in combination with partitioning; the amount of resources an entity consumes is limited by assigning each entity an ‘area of interest’. There are many ways to perform interest management, both structured and unstructured, and several challenges remain (the reader is guided to Yahyavi and Kemme [23] for details).

Latency can be defined as “the delay between execution of an update at the primary copy of an object and the replica receiving the object update” [23]. One of the critical aspects of MMOWs is their real-time **responsiveness**, which demands “message latency should be minimized while bandwidth use should be efficient” [6]. From a human perspective, real-time responsiveness is achieved when “the time between the event being generated and the time it is executed and perceived by the user is unperceivable [sic]” [13] *i.e.*, responsiveness despite latency. The tolerance threshold for latency in games is between 100 and 300 milliseconds, depending on the game type (ranging from FPS games to RPGs) [23]. Liu, Bowman, and Chang [13] report a trade-off between consistency and responsiveness (throughput).

If multiple servers are collaborating to maintain a compute space, then **consistency** between the state of each server must be maintained despite networking delay. Yahyavi and Kemme [23] report a well-known trade-off between performance (availability) and consistency restrictions. Consistency involves having a primary copy of each object and sending updates (*i.e.*, update dissemination) to replicas, in such a way that the causal order of events are consistent [23, 13]. Note, the difference between objects being replicated here, and entire server partitions being replicated as with shards. A way to assess a degree of inconsistency is by comparing the (potentially inconsistent) state of each replica against a virtual perfect replica [23].

Considering the prevalence of smartphones, some MMOWs have moved a mobile platform. But the limitations in networking and computing power of such devices are a factor [23]. Taking into account whether to use local or remote resources, is important for efficient **resource utilization**.

For MMOWs using a centralized cluster of servers, the **availability** of end nodes in the network is not so much an issue, but in P2P, nodes are “much more prone to failures or unscheduled disconnections” [23], which can adversely affect the network topology [6]. Related to the availability of the network, nodes must be fault tolerant and support data persistence for recoverability. Data persistence means the ability to save and access world states despite disconnections, which remains a challenge in P2P systems [6]. Having re-

dundant backup copies of primary objects, redundant network connections and redundant servers [5] are ways to provide fault tolerance.

In games, one of the main concerns is cheating, which corresponds to **security** in other applications. Yahyavi and Kemme [23] present three categories for cheating: Interrupting Information Dissemination, which includes premature disconnection, flooding of the network, replay attacks and the dropping of updates to peers; Illegal Game Actions, which includes tampering of end nodes of the network, falsifying identity and the use of computer enhanced data where human readings are expected; and Unauthorized Information Access, which includes the tampering of end nodes or network traffic analysis in order to gain access to privileged information. It is easier to build a game using a centralized architecture; if a P2P communication is used instead of a centralized approach, then dealing with cheating (security) and maintaining control over the game remains a challenge [23, 6]. Having more control means that the architecture is easier to manage and maintain [23].

3 Evaluating Properties Against the Domain of IoT

IoT can be discussed from two perspectives, ‘Internet’ centric and ‘Thing’ centric [4, 1]; leading to different architectures *i.e.*, centered on cloud computing (‘remote’), or centered on the user (‘local’), respectively [4]. A complication that MMOWs do not currently have to deal with is that IoT network architectures can be (partially) disconnected *e.g.*, MANets [14]. If (partially) disconnect networks are present, interaction with the local environment is most likely still possible, contrary to remote access.

Because of the projected size of the IoT, **scalability** is expected to be a major issue and often mentioned in literature [1, 4, 14, 15]. Literature is divided on whether a centralized cloud computing, decentralized P2P, or a hybrid architecture will achieve the needed scalability (through the addition of resources) for IoT. Arguments for cloud computing are based on the advantages of abundant cloud storage and processing capabilities, while being tightly controlled for efficient energy usage and reliability [4]. Counter arguments state that decentralized P2P architectures show promise [4, 1] and that a centralized architecture cannot lead to a truly scalable solution [14] *i.e.*, wanting to reap the “seemingly endless amount of distributed computing resources and storage owned by various owners” [4]. If billions of devices will be soon connected to the Internet, producing a potentially quadratic number of interactions [23], then somehow IoT needs to facilitate those interactions [19]. Currently, many sensor network solutions are sense-only solutions, with many-to-one interactions *i.e.*, only a small fraction of solutions are sense-and-react applications (using actuators), with many-to-many interactions [16]. To achieve scalability through reducing consumption, there seems to be only one construct mentioned in IoT literature. From the domain of WSNs, the construct of ‘clustering’, for both physical devices and also objects present on those devices, has been suggested for scalability; either abstract regions are formed [16] or a few nodes are elected as ‘cluster heads’ to act as decentralized authorities for each cluster [14]. A problem with current clustering protocols being that mobility of nodes has hardly been considered [14].

To achieve the appropriate **responsiveness**, IoT is said to require two classes of traffic: the throughput and delay tolerant elastic traffic; and, the bandwidth and delay sensitive inelastic (real-time) traffic. These two types can be further discriminated into various levels of quality of service [4]. Perera et al. [19] refer to two different classes of traffic, namely ‘event driven’ versus ‘time driven’, which correlates to ‘event-triggered’ and ‘periodic’, respectively, in WSN literature [16]. Perera et al. [19] mention that “real-time data processing is essential”.

“The synchronisation [sic] of data across the architecture” (*i.e.*, **consistency**) is envisioned to be a big challenge in IoT [14]. In WSN literature, four approaches emerge on how to provide access to data: similar to a relational database, as remotely accessible

variables or tuples, through mobile code, or by message passing [16]. A number of projects in IoT use ‘participatory sensing’, where sensors on people are used to obtain readings local to the user [4]. A difficulty in IoT is that (partially) disconnected networks exist, so that “there is a potential for various non-homogeneous copies of object data across the architecture”, making support for ‘one-copy’ [primary copy] equivalence problematic *i.e.*, “the value of all copies should be identical after a transaction” [14].

The aim of IoT is to include many platforms and devices with various resource limitations *e.g.*, cloud server clusters, desktop computers, mobile devices, sensor networks and everyday objects; access to resources is also through various network infrastructure, such as a wireless personal area network (*e.g.*, Bluetooth), wireless local area network (*e.g.*, Wi-Fi), wireless wide area network (*e.g.*, 2G and 3G networks), and satellite network (*e.g.*, GPS) [19]. Not all ‘things’ in the IoT have sufficient energy and processing power to do comprehensive data processing (*e.g.*, sensors in a WSN); in that case, data must be networked to more powerful devices and processed there [19]. Because billions of IoT devices will potentially be communicating with each other, IoT research has noted the importance of optimizing computation in the various ‘layers’ of the IoT infrastructure [19] *i.e.*, **resource utilization**. An example of this would be the “notion of distributing computation in order to reduce the communication overhead, which is generally termed in-network processing or in-network computing” [4]. In participatory sensory, people are centric [4], but this can be generalized in the sense that interactions for an object in the IoT is highly dependent on the objects surroundings *e.g.*, presence of other objects or people [19].

With the existence of partially or permanent disconnected networks in IoT [14], robustness and fault tolerance (*i.e.*, **availability**) will become fundamental research topics in IoT [15]. Considering the ‘extremely large scale’ of IoT and the ‘high level of dynamism in the network’, self-organization is suggested as a solution [15] *e.g.*, supporting merges and splits of the network [14]. In WSN literature, very little research has been done with respect to having mobile nodes in a network [16]. And, programming approaches for WSNs provide for only limited guarantees in the face of the various types of hardware faults [16].

Perera et al. [19] state that the “IoT paradigm will intensify the challenges in **security** and privacy”. Security is related to concepts such as authentication, privacy and integrity. For IoT, cryptographic algorithms commonly used in authentication are typically problematic in devices with various resource limitations (*e.g.*, sensor networks), using large amounts of energy and bandwidth [1]. Authentication usually involves identifying people, but in IoT identities are associated with objects [15]. If authentication is to be possible in (partially) disconnected environments, the procedure will have to be possible locally [14]. Privacy refers to the access of data related to an individual [15]. While surfing the Internet, individuals usually play an active role in their privacy, but in IoT sensors are expected to collect information about individuals passively, without them actively using an IoT service and without control over what information is being collected [1, 19]. Once the information is generated, it will most like be retained indefinitely, unless a mechanism is in place to allow for ‘digital forgetting’ [1]. The purpose of data integrity is to prevent adverse modification of data without detection [1, 14]. A criticism with WSNs and thus IoT, is that hardware components are easy to physically attack, because they are largely unattended and that if wireless communication is used, it can be eavesdropped [1]. Security and privacy are very much open issues in IoT [19, 15].

4 MMOWs and IoT, Comparison of Key Properties

To achieve **scalability** through adding resources, cloud computing and P2P solutions are being considered in the domain of IoT, but research in the domain of MMOWs is more advanced, with running platforms utilizing partitioning schemes such as regionalization and

replication in combination with interest management. To reduce consumption, clustering has been suggested in IoT, for devices and objects present on those devices. Clustering for devices in IoT can be equated to the partitioning techniques regions and shards in MMOWs; and, clustering for objects in IoT can be equated to interest management in MMOWs [23, 14]. For example, López et al. [14] suggests using clusters determined through the use of context information, which is very similar to the existing MMOW project called *Donnybrook* using ‘interest sets’ [23]. There seems to be an overarching opinion that a single architecture or middleware will enable IoT, but there is a risk of fragmentation [15]. With many working on their own platform, one can speculate if IoT will be a collection of clouds, also facing interoperability problems found in multicloud computing *e.g.*, the locking in of clients [21]. If IoT does fragment and considering the advances in the domain of MMOWs, this means solutions from the domain of MMOWs (specifically the Metaverse; see Dionisio, Burns III, and Gilbert [3]) could be exapted to the domain of IoT.

In the IoT literature consulted, there was no mention of the consistency and **responsiveness** trade-off found in MMOWs. Real-time data is one of the properties that “distinguish virtual worlds [MMOWs] from other distributed systems” [13]. Real-time traffic will not be required in all applications of IoT, but those applications that require delay sensitive inelastic traffic would be in the same class as an MMOW. Contrary to cloud-based MMOWs, response times in IoT can be kept lower, if interactions are kept in a local environment rather than with remote resources, due to reduced networking latency.

Also not in IoT literature, was any mention of the performance (availability) and consistency restrictions, in a partitioned space *i.e.*, a distributed system. A metric for **consistency** in MMOWs is ‘drift distance’ [6]; for a moving entity in a virtual world, the drift distance is, the absolute value of: the distance between the position of the entity locally and its position remotely. When dealing with the physical world, the drift distance could be calculated between the local position of the entity and its physical position. This means that given a physical entity with multiple virtual replicas (*i.e.*, multiple versions of reality) the copy most similar to local physical entity is the one most consistent.

Until rather recently MMOW architectures have had the luxury of only having to support the desktop class of devices, connected to a server cluster, with MMOWs on mobile devices only recently becoming more common. In IoT, efficient **resource utilization** is still very much a challenge. Some sensors are able to produce a continuous stream of data and networking that data to more powerful devices or the cloud might outweigh the benefit, given the limited computing power of sensors and mobile devices. If faced with (partially) disconnected networks, the benefit of accessing local resources must be considered. If MMOWs are to interface with IoT, this means MMOWs must also face (partially) disconnected networks. To allow for the use of local resources in a disconnected state, authority over local resources can be delegated to the local environment; a caveat being that, although MMOWs have considered P2P or hybrid architectures, engines have typically been centralized clusters, rather than geographically distributed systems.

Similar to resource utilization, IoT research has already taken into account (partially) disconnected networks, but if MMOWs are to interface with IoT, issues with **availability** must be dealt with. If a shard is disconnected, delay tolerant networking must be used so that the shard can be merge when a connection is established [18, 2].

Authentication is not mentioned as an issue for MMOWs as long as a central authority is used. If decentralized P2P is used for either IoT or an MMOW, **security** remains an issue. Miorandi et al. [15] state that IoT needs to move away from centralized approaches, to a fully distributed and dynamic approach. Privacy in an MMOW is similar to authentication, except that rather than trying to access private player data, attackers are trying to gain sensitive information pertaining to game state. Most game cheats are primarily against data integrity. Similar to IoT, end nodes in an MMOW cloud are also easily attacked and the networking tampered with. Premature disconnections and decentralized P2P for MMOWs has a likeness to partially disconnected networks in IoT. Missing from

the IoT discourse, is the issue of maintenance; if IoT is to be enabled through a single platform, how will updates to the platform be managed?

5 Case Studies of an MMOW Interfacing With or Mediating IoT

In previous sections, properties have been found in the domain of MMOWs, that overlap with the domain of IoT. To show that these two domains do indeed overlap, three different case studies are presented that combine MMOWs with IoT. The case studies include: (i) MediaSense, a middleware to handle communication between nodes of a distributed IoT network; (ii) Immersive Networking, a framework based on the MPEG-V standard, which outlines an architecture for interoperability between virtual worlds (MMOWs) and the physical world (via IoT); and (iii) the Metaverse as an aggregate of services from information systems, including IoT.

5.1 MMOW Interfacing with a Distributed IoT Middleware

MediaSense [12] is a distributed IoT platform based on the Distributed Context eXchange Protocol (DCXP) for sharing context information between peers; locations of peers are denoted by Universal Context IDs and resolved via a Distributed Hash Table. Each end point in the middleware is referred to as a Context User Agent, which stores context information in the Distributed Hash Table as a $\{\text{key}, \text{value}\}$ pair. Context information is modeled as relations between objects, which are maintained via DCXP methods (GET/SET or PUBLISH/SUBSCRIBE). The platform (see Figure 1) allows for a direct mapping of events between the virtual objects of a MMOW by: (1) collocating a Context User Agent (running in a ‘bootstrap node’ [11] of the DCXP P2P context network) with a node from the MMOW cloud; (2) creating an interface between collocated nodes so that virtual objects can be assigned to the Universal Context IDs of corresponding physical devices, on which a Context User Agent is running; and (3) mapping game events (between virtual objects) to DCXP methods (between Context User Agents). New physical nodes can announce themselves and claim to represent a MMOW virtual object. Any node in the DCXP P2P context network can participate in the MMOW cloud; the only requirement being that at least one DCXP Context User Agent is collocated with the MMOW Cloud. Extrapolating the scenario means all objects are hosted in DCXP nodes, so that the game world (consisting objects and relations) is hosted by the DCXP P2P context network *i.e.*, effectively distributing the game engine. To achieve scalability for physical nodes, while maintaining relationships with other objects, Walters [22] presents a method for clustering objects according to multi-criteria ranking.

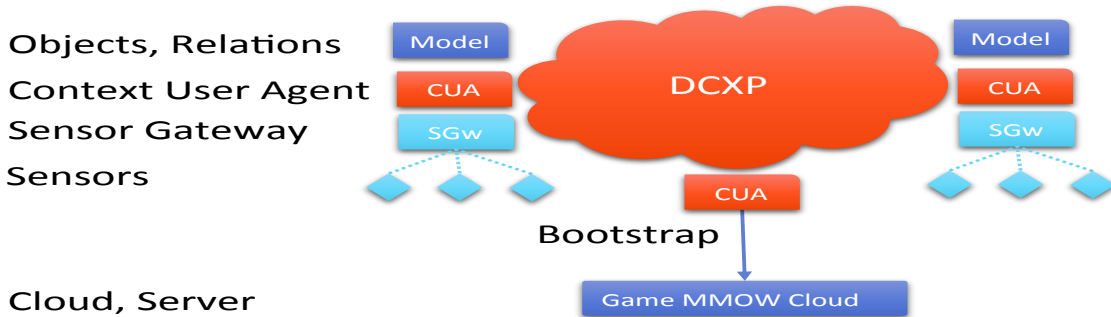


Figure 1: Connecting MMOW and a DCXP P2P context network

5.2 IoT Interfacing with Virtual Environments on a MMOW infrastructure

MPEG-V [7] is a standard which outlines an architecture for interoperability between a virtual world (MMOWs) and the physical world (via IoT). In addition the standard aims to enable interoperability between virtual worlds by characterizing metadata for virtual world objects, so that objects can be migrated from one virtual world to another. MPEG-V has been drafted to support the Metaverse *i.e.*, forming a system of interoperable interconnected virtual worlds. Part 1 of the MPEG-V standard outlines a top level architecture, with target areas for standardization.

Immersive Networking [10] aims to go beyond IoT to enable the massive and scalable sharing of context information across a MMOW infrastructure *i.e.*, information sharing in real-time for timely and intelligent decisions in different user scenarios. Immersive Networking is a framework technology that allows users to share context information in multiple practical scenarios to obtain immediate, objective feedback from the pervasive game engine towards a relational IoT solution which is identified as interfacing with a virtual environment on a MMOW infrastructure. The Immersive Networking framework is envisioned as a MPEG-V virtual world infrastructure with a distributed controller for interfacing with virtual environments, as depicted in Figure 2. The framework specifies the context and semantics required to provide interoperability in the distributed controlling of virtual agents as well as generic virtual objects. The framework aims to replace the adaptation engine (RV or VR) in MPEG-V, making seamless experiences in virtual environments possible through the self-organization of connectivity between people, places and things.

5.3 Virtual Worlds as a Behind the Scenes Resource

Nevelsteen [17] concluded that a virtual world engine is in the same product line as a game engine for pervasive games. Considering pervasive games need a sensory system to monitor the physical world (through the use of non-standard input devices), IoT could potentially serve as such a sensory system. Since a MMOW is a virtual world at a massive scale, this means that a MMOW can be exploited as a ‘behind the scenes’ resource for coordinating and managing devices and interaction in the physical space.

Furthermore, according to Rehm et al. [20], the concept of IoT has been recently been replaced by the concept of Cyber-Physical Systems. Rehm et al. believe that “virtual

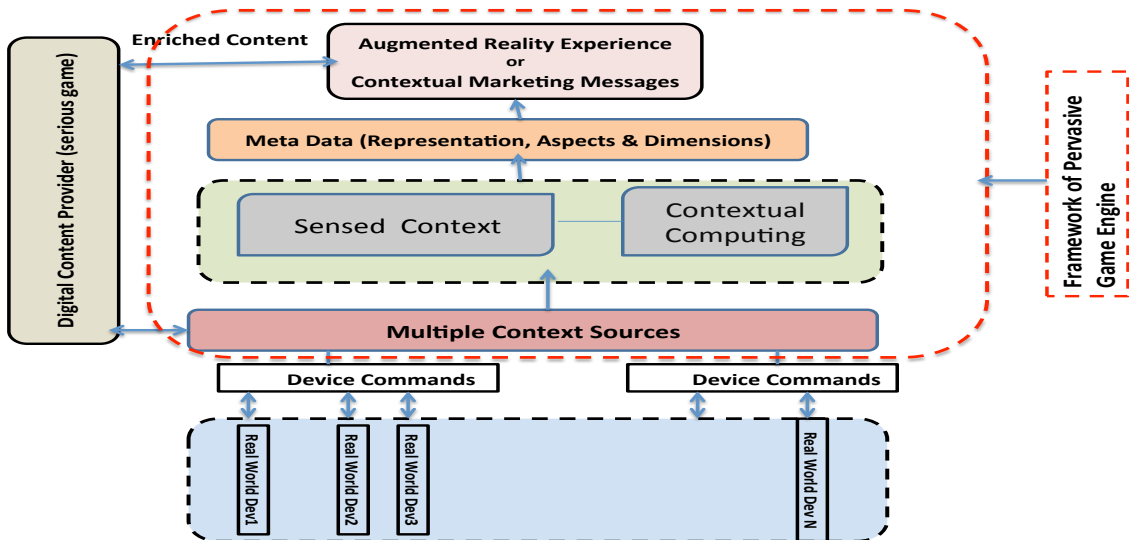


Figure 2: **Immersive Networking framework** - framework that aims to replace the adaptation engine (RV or VR) in MPEG-V standard

worlds can serve as platforms to facilitate the integration required by CPS [Cyber-Physical Systems]”. And, by extrapolation, they “conceive of a unified platform, the Metaverse, built on VW [virtual world] technologies that allow for the integration of technological, physical, and human elements of CPS [Cyber-Physical Systems]”.

6 Conclusions

The result of this article is the “explication of the problem” [9] of scaling IoT, and incorporating a MMOW with IoT into a pervasive system. Six properties, specific to a massive number of entities interacting, have been identified; first in the domain of MMOWs, second in the domain of IoT, and then compared in discussion. IoT can clearly learn from advances, in availability with respect to P2P systems, and scalability, from the domain of MMOWs. When virtual worlds start to incorporate IoT, blending the virtual and the physical, MMOWs can clearly learn from advances made, in resource utilization, availability, and responsiveness with respect to (partially) disconnected networks, from the domain of IoT. For consistency and security, there seems to be advances in both domains that can cross over to the other domain. Rather than think that IoT brings about a whole new set of issues, advances in other domains (*e.g.*, the Metaverse) should be considered during IoT development.

Author’s contributions

Research and drafting of property requirements relating Massive Multiplayer Online Worlds and the Internet of Things is done by Nevelsteen. Conceptual work on decentralizing a game engine for use with Internet of Things is done in collaboration with all three authors. Each of the case studies is provided by a single author; MediaSense by Kanter, Immersive Networking by Rahmani and Virtual Worlds as a “behind the scenes” resource by Nevelsteen.

Acknowledgements

Research was made possible by a grant from the Swedish Governmental Agency for Innovation Systems to the Mobile Life Vinn Excellence Center.

REFERENCES

- [1] Luigi Atzori, Antonio Iera, and Giacomo Morabito. “The Internet of Things: A survey.” In: *Computer networks* 54.15 (2010), pp. 2787–2805. DOI: 10.1016/j.comnet.2010.05.010.
- [2] Isabelle Demeure et al. “Transhumance: a Platform on a mobile Ad hoc NETwork challenging collaborative gaming.” In: *Collaborative Technologies and Systems, 2008. CTS 2008. International Symposium on*. Irvine, CA, USA: IEEE, 2008, pp. 221–228. DOI: 10.1109/CTS.2008.4543935.
- [3] John David n. Dionisio, William G. Burns III, and Richard Gilbert. “3D Virtual Worlds and the Metaverse: Current Status and Future Possibilities.” In: *ACM Computing Surveys* 45.3 (2013), 34:1 –34:38. DOI: 10.1145/2480741.2480751.
- [4] Jayavardhana Gubbi et al. “Internet of Things (IoT): A vision, architectural elements, and future directions.” In: *Future Generation Computer Systems* 29.7 (2013), pp. 1645–1660. DOI: 10.1016/j.future.2013.01.010.

- [5] Thorsten Hampel, Thomas Bopp, and Robert Hinn. “A peer-to-peer architecture for massive multiplayer online games.” In: *Proceedings of 5th ACM SIGCOMM workshop on Network and system support for games*. 48. New York, NY, USA: ACM, 2006. DOI: 10.1145/1230040.1230058.
- [6] Shun-Yun Hu, HuJui-Fa Chen, and Tsu-Han Chen. “VON: a scalable peer-to-peer network for virtual environments.” In: *Network, IEEE* 20.4 (2006), pp. 22–31. DOI: 10.1109/MNET.2006.1668400.
- [7] ISO/IEC JTC 1/SC 29. *ISO/IEC 23005, Media Context and Control (MPEG-V)*. Tech. rep. International Organization for Standardization (ISO), 2009.
- [8] ISO/IEC JTC 1/SC 7. *ISO/IEC 25010:2011 Systems and software engineering – Systems and software Quality Requirements and Evaluation (SQuaRE) – System and software quality models*. Tech. rep. International Organization for Standardization (ISO), 2011.
- [9] Paul Johannesson and Erik Perjons. *An Introduction to Design Science*. Switzerland: Springer International Publishing, 2014. DOI: 10.1007/978-3-319-10632-8.
- [10] Theo Kanter, Uno Fors, and Rahim Rahmani. “Immersive Networking – A Framework for Virtual Environments with Augmented Reality in Human Decision-Making.” In: *International Journal of Multimedia and Ubiquitous Engineering* (in press).
- [11] Theo Kanter et al. “Distributed context support for ubiquitous mobile awareness services.” In: *Fourth International Conference on Communications and Networking in China, 2009. ChinaCOM 2009*. IEEE. Xian, 2009, pp. 1–5. DOI: 10.1109/CHINACOM.2009.5339728.
- [12] Theo Kanter et al. “MediaSense – an Internet of Things Platform for Scalable and Decentralized Context Sharing and Control.” In: *ICDT 2012*. Chamonix, France: IARIA, 2012, pp. 27 –32. ISBN: 978-1-61208-193-9.
- [13] Huaiyu Liu, Mic Bowman, and Francis Chang. “Survey of state melding in virtual worlds.” In: *ACM Computing Surveys (CSUR)* 44.4 (2012), p. 21. DOI: 10.1145/2333112.2333116.
- [14] Tomás Sánchez López et al. “Architecting the Internet of Things.” In: Berlin Heidelberg: Springer-Verlag Berlin Heidelberg, 2011. Chap. Resource Management in the Internet of Things: Clustering, Synchronisation and Software Agents, pp. 159–193. DOI: 10.1007/978-3-642-19157-2_7.
- [15] Daniele Miorandi et al. “Internet of Things: Vision, applications and research challenges.” In: *Ad Hoc Networks* 10.7 (2012), pp. 1497–1516. DOI: 10.1016/j.adhoc.2012.02.016.
- [16] Luca Mottola and Gian Pietro Picco. “Programming wireless sensor networks: Fundamental concepts and state of the art.” In: *ACM Computing Surveys (CSUR)* 43.3 (2011), p. 19. DOI: 10.1145/1922649.1922656.
- [17] Kim J. L. Nevelsteen. *A Survey of Characteristic Engine Features for Technology-Sustained Pervasive Games*. SpringerBriefs in Computer Science. Switzerland: Springer International Publishing, May 2015. DOI: 10.1007/978-3-319-17632-1.
- [18] Kim J. L. Nevelsteen. “‘Virtual World’, Defined from a Technological Perspective, and Applied to Video Games, Mixed Reality and the Metaverse.” (in press).
- [19] Charith Perera et al. “Context aware computing for the Internet of Things: A survey.” In: *Communications Surveys & Tutorials, IEEE* 16.1 (2014), pp. 414–454. DOI: 10.1109/SURV.2013.042313.00197.
- [20] Sven-Volker Rehm, Lakshmi Goel, and Mattia Crespi. “The Metaverse as Mediator between Technology, Trends, and the Digital Transformation of Society and Business.” In: *Journal For Virtual Worlds Research* 8.2 (2015). DOI: 10.4101/jvwr.v8i2.7149.
- [21] Mukesh Singhal et al. “Collaboration in Multicloud Computing Environments: Framework and Security Issues.” In: *Computer* 46.2 (2013), pp. 76–84. DOI: 10.1109/MC.2013.46.
- [22] Jamie Walters. “Distributed Immersive Participation – Realising Multi-Criteria Context-Centric Relationships on an Internet of Things.” 1. Department of Computer and System Sciences, Stockholm University, 2014. ISBN: 978-91-7447-987-4.

- [23] Amir Yahyavi and Bettina Kemme. “Peer-to-peer architectures for massively multiplayer online games: A survey.” In: *ACM Computing Surveys (CSUR)* 46.1 (2013), p. 9. DOI: 10.1145/2522968.2522977.